- 5-7 pages.

EDITED TRANSCRIPT

D. W. DAVIES interviewed by M. CAMPBELL-KELLY
at the National Physical Laboratory, UK
Monday, 17 March 1986

MCK I am here at the National Physical Laboratory, the date is the 17th March 1986, the time is 3:40, and I am talking to Donald Davies formerly Superintendent of the [Computer Science] Division. Donald, in this interview I want to focus on your data communications work with the National Physical Laboratory, which took place over the period 1965-1975. But first, tell me briefly about your educational background, the work you did with the "Tube Alloys" project during the war, and how you came to be appointed to the NPL.

DWD Well, I started physics first of all at Imperial College and graduated in 1943, with a first class honours degree in physics, and at the time you were immediately drafted to a wartime project. It was [C. P.] Snow who was the chairman of the committee who did this allocation, and I was sent off to Birmingham University to work on a "Tube Alloys" project. Now that was interesting in a way because it involved a lot of numerical calculations and I spent most of my time supervising groups of computers, who were people with Brunsvigas and electric calculators. So obviously the need for more rapid and efficient calculation was ground into me during that time. The work was mostly on the design of the [U235] diffusion separator plant. (Incidentally my supervisor for part of the time was the famous traitor Klaus Fuchs, who I grew to know quite well.) I went back to Imperial College after the war to take a degree in mathematics, which you could do with comparatively little extra effort. It was an applied mathematics type of degree. That's also significant in that Hyman Levy taught numerical analysis and inspired people with a real interest in numerical analysis which was not much studied before the war in universities - it was treated as a rather ordinary sort of subject. He was really very interesting and so I had some further involvement in numerical work at that time. During that time [Norbert] Weiner came to lecture there and I went to his lecture and was very inspired by it. But it wasn't of course about digital computers, but about the concepts he called 'cybernetics'. He was certainly an inspiring speaker, although when I heard him a few years later give exactly the same

lecture, I began to wonder if the real promise of that work was as much as it seemed. Strangely enough, I also met him while I was at MIT on a fellowship and at that time I'm afraid he appeared to be in his dotage. A lot of things certainly inspired me to take an interest, but also the head of the maths division here, [John} Womersley, came down and gave a talk. After the talk I realized that there was an opportunity to get into this very interesting business, and went up to him at the lecture and said "Well, how can I get to NPL?", and I came down here and was interviewed and so on, and that's how I got to NPL.

MCK At the NPL you were involved particularly with the Pilot ACE and the full-scale ACE computers which were designed between 1945 and the late-1950s, of which a lot has been written. But in 1960 you were promoted to Deputy Superintendent of the Division of Autonomics. So please take a few minutes to tell me about the work of the division in the early 1960s. What was the organizational structure, how many research groups were there, and what were your particular interests?

DWD Yes, well ... it's a long way back. I think there were about four research groups and Uttley [A.M. Uttley, the Superintendent] built them around the interests of the people. He came round to see the various people who he thought could lead research and asked them what they were interested in. I particularly remember that Ted Newman was given charge of pattern recognition and built up quite a large team in both visual and voice pattern recognition. I expressed an interest in the mechanisms of computing and high speed computing elements and so on, and began to work on the cryotron, which was at that time my main interest. But then as you probably know I got more involved with the Advanced Computer Technology Project. The other ones I'm rather pushed to actually remember, I'm sorry...

MCK How many people were there, about, in the division?

DWD I have an impression that there were about 60 when Uttley began, and this increased, I think, to close on 100. This was at a time when the idea of the NPL was to become a kind of University of Teddington and was very different to what we have now, in fact almost the reverse. Of course 'Jttley's main interests were in neuro-physiology and the analogy between the brain and the computer and so on, so I should have obviously remembered that one of the big

interests here was in the physiology of the brain, and there was a physiological group. There was also an adaptive control group under [Percy H.] Hammond. I think I've probably therefore just about missed out one group, but I can't remember. But essentially Uttley had the idea of building the research teams around the interest of the people, quite unlike the thing which followed which was the Requirements Boards. There was little thought of immediate outcome from any of the work - very much long term work.

MCK One of your jobs was to administer the Advanced Computer Technology Project. What were the origins of the project, how was it organised, and what was your particular role?

DWD Again, I have to search back in my memory. I think it was the announcement that very large computers were going to be built in USA. I think it was the announcement of the STRETCH Project, which of course took a long time to develop but began quite early. Harwell called a big meeting of all those people interested in computing at the time - it was a meeting of about fifty people - and it was decided at the end of that meeting that there were two actions that should be taken. One was to somehow sponsor a large project, using existing technology, to build the biggest computer we could have; that got turned into the Atlas project. The other one was to make sure we could advance our own technology in Britain, so that we were going to have our own technology for the next wave of computing. Now at the time the deputy director [of Mintech] Ieun Maddock was given the job of organizing this second part of the thing and we had an enormous number of meetings with people at the newly formed Ministry of Technology. It was the beginning of [Prime Minister Harold Wilson's] "white-heat" of technical revolution and so on, and after a great deal of argument the terms that the treasury were going to allow for research were agreed. I think that may have been the first time the treasury laid down its fifty per cent rule - that industry must put in fifty per cent. I was appointed the project leader, and I think it was intended really that it should be a full time role, but in fact the contract started very slowly because of difficulties in getting the contract terms approved and the first contract was much delayed - very much like Alvey. Although the amount of money was quite large for the time - it was £1 million per year, which if scaled up is not that incomparable with Alvey - it was actually run by a couple of people at headquarters and myself with an assistant here, and the assistance of anybody else in the laboratory on whom I wanted to call. So really there was

a lot of part time effort at the NPL helping me. It was my job to go out and see the companies and find out what they wanted to do,to get the terms of the contract written (that is the technical aspects of the contract), to arrange such things as the terms of the levies and get the contractual problems to sort out. Actually ACTP had quite a number of successes. Eventually it was closed down by government edict, but among the successes was the Distributed Array Processor [DAP] and the Content Addressable File Store [CAFS]. They took an enormous time to get developed to a saleable state, which I think was due to the lack of interest of ICL in these projects. Once they had passed the research stage there was an enormous gap before they turned into development projects. It was, I think, in its day a worthwhile operation.

MCK Your interest in computer communications seems to date from your trip to the States in April 1965, when you attended the IFIP '65 Congress and made a number of research visits. I know you visited places where cryogenic storage research was being done, in which you had a research interest. But you also visited several of the early time-sharing groups, such as Project MAC and IBM. What was the reason for those visits, and how did the work compare with what was going on in England?

DWD Starting at the beginning, actually its not really true that my interest in computer communications dated from then. But that was the point at which I became sufficiently enthused that I was willing to put a lot of the NPL's effort into it, and I was able to do so of course when I became Superintendent. In fact I'd taken an interest in telephone switching right from childhood. I used to design relay circuits and so on. (I had no chance of putting them into operation; they were complicated logic based on relay circuits, which was the only way one could design logic in those days.) Also, quite a lot earlier, I don't remember the date, I got involved in a classified project to send information securely. Even now, I shouldn't say where it was from and to, but it was from a testing range back to the place where the data was going to be massaged or computed with. This involved both data communications, inventing protocols, and data security; it was actually very primitive, but it was an early bit of experience in that field. They used telex lines - that shows you how slow it was.

Now, the reason for the visits associated with the IFIP Congress. Well, I was by that time in the habit of going to Congresses in USA - somehow I was able to work up the case for doing so - and it was generally thought that

whenever you went to USA, you ought to justify the expenditure by making some visits. So there was in some ways a pressure on you to make a number of visits anyway, and I was also aware of the enthusiasm for time-sharing but actually I didn't share it a great deal at the time that I went. But I took the opportunity to go and visit these places where time-sharing was done. It was not so much Project MAC that impressed me as the JOSS Project at the RAND Corporation, and also the BASIC project at Dartmouth, because they seemed to me really aimed at making computing useful by the average person. Project MAC seemed to concentrate very much more on the technology, it also went back to something that I experienced earlier when I had a Commonwealth Fund Fellowship to study at MIT for a year in 1954. While I was there, I was impressed by the fact that, as compared with our computer on which everybody was rushing around doing real jobs like calculating aircraft structures and things of that sort, at MIT (apart from the classified projects which I wasn't supposed to know about - they were on a different floor) all the people worked only on operating systems, and nobody actually used the machine. There were these marvelous operating systems, but nobody was rushing around to actually use them. At MIT there was much more interest in the technology of how you made a computer available than in actually using it. Project MAC clearly was an exception and it became a very important project for the user as well. But that's the background to it - it wasn't that I was desperately interested in time-sharing. It was one of the things that I did while I was there to look up projects in cryotron research.

Now, as for how the work compared with what was going on in England. Time-sharing in Britain was already known, but it meant something different, it was usually the way in which the processor shared its time between the calculation and the I/O of its programs. Multiprogramming is probably a better term for it. It wasn't making it available to lots of users. As far as I can make out, there weren't any projects at that time (although I might have been missing something in a university somewhere). It was really because of that that Brian Shackle [see later] mentioned the idea on the way back of organizing this meeting in Britain to let more people know about it. I think, in fact, a lot more people knew about time-sharing, because there was a great deal of publicity coming across the Atlantic about Project MAC, probably more than was justified - because compared with the other time-sharing projects in USA everybody seemed to be talking about Project MAC. Now, the meeting in Britain really was organized in the form of two meetings [in November 1965]. First of all we had a kind of specialist seminar at NPL

lasting three days. And then, I think on the following day, we had this open meeting which was much more a popularizing meeting - it was to tell people how it worked rather than go into details about the design. That was run by the British Computer Society who could attract a much bigger audience - they might be able to give you more information about who spoke at that meeting. Among the people who came from MIT, I have the impression that there was something like between eight and ten people.

MCK Who funded them?

DWD I don't know. I have forgotten now how it was done. Somehow I think it would be unlikely that we would get the funds, but we might have found some way of doing it. I'm afraid I've forgotten that.

MCK Shackle was with EMI?

DWD Yes, that's right, he was with EMI. He was doing research in human factors then, I think; he had some government projects, how they were funded I just don't remember. Among those ten people were [Jack B.] Dennis, [Fernando J.] Corbato who was much involved in the TSS, and Larry Roberts of course. A chap called Richard Mills [of MIT] who I think was one of the people who had better ideas than anyone on the communications aspects. I'd spoken to him while I was at MIT on my visit. Larry Roberts didn't take a very big part in the discussions. Actually, most of the discussions tended to be about the operating system aspects, but certainly the mismatch between time-sharing and the telephone network was mentioned. It was that which sort of triggered off my thoughts, and it was in the evenings during that meeting that I first began to think about packet-switching.

MCK Well, in March 1966 you gave a seminar on packet-switching to an invited audience at NPL, and the original time-sharing seminars had taken place in November 1965; so somewhere between those dates you'd invented the packet-switching concept. Tell me exactly how the ideas evolved, it sounds as though they came in a flash during those early first seminars?

DWD Yes, I think that's virtually true. The basic ideas were produced really just in a few evenings of thought, during or immediately after the seminar. Then I began to work them out in a little bit more detail, thinking about the

interfaces involved, about how the processes would be organized and drawing a few diagrams and so on. But the initial idea was quite a simple one and didn't take more than an evening or so to work out in basic form. So I think it took about two months to develop to the point where I felt it was worth showing the ideas to people and handing the papers to them. And then the main event, as you said, was the lecture in March 1966. After that lecture, I was still concerned whether the ideas had any merit: were they very obvious, and did everybody know about them anyway? And the relationship to message-switching obviously interested me a great deal.

MCK How much did you know about message-switching at the time; you've mentioned that you had a bit of a background in telecommunications, but we didn't have much in the line of message-switches in Britain at that time.

DWD Well, certainly I knew about message-switching, I don't know whether there were many in this country, but I knew about this and had a little book about the way it operated. Of course one of the first thoughts was "well, this is just message-switching by another name".

MCK Was that in an amateur capacity?

DWD Yes. Certainly I didn't have any need for any professional interest in message-switching, but it is one of the things you read in the journals and so on. Yes. I'd never had any practical experience of it, but later on of course I went round and saw message-switching in operation. As you know, the main differences in the way message-switches work are the lower speeds, and they don't multiplex in quite the same way. They don't have very short messages and they don't concentrate on reducing the store-and-forward delay. Also, because they take responsibility for the message, they lay a lot of emphasis on never losing a message and consequently tend to use magnetic stores, which again slows the thing down. Also they are prepared to hold messages indefinitely, almost, if the receiver doesn't happen to be ready to receive them; so they also act as fairly large buffers. Certainly I knew about message-switching at the time, yes.

MCK Now for the 1966 seminar you invited several people from the Post Office. I'm not quite clear what your relationships with the Post Office were at that time? Were they old friends?

DWD Well, obviously I knew a number of people. I visited Dollis Hill Research Station on many occasions in the early days of ACE, and also we used the Post Office as a source of equipment at times, so I had relationships in different places throughout the Post Office. I don't think I had any contacts at a high level in the Post Office at that particular time, though I'm not quite sure about that. In fact, I don't recall that I actually invited anybody by name. It's very likely that they came as a result of a general invitation going out on the notice board. I don't think we took any special precautions to invite lots of people, which is rather surprising; it gave us a tremendous surprise when about a hundred and twenty people turned up and overfilled the lecture theatre. It was something really very unusual to have people standing at the back and so on, so this immediately told me that the subject was one of great interest, and also the fact that the Post Office people listened with respect and didn't immediately say "this is a load of rubbish", gave me enormous encouragement. The only comments in discussion I remember was "This is just the same thing as message-switching", and of course I had a reasonable answer to that. A lot of people said we will never convince anybody in the Post Office because they were so fixed on the idea of circuit-switching, and that indeed proved to be true.

Now, the relationship with the Post Office is an interesting question because right from the beginning I could see the danger that the Post Office would say "Oh, yes, we are very interested in this, we will set up a team of two or three people to work on it, and of course you, NPL, don't need to, therefore". And all they had to do was speak to the director and he would say "OK, very interesting, we've cottoned on to this idea but its really the Post Office's business" and I would have nothing further to do with it. That was the thing I most feared, so I wasn't all that anxious to go right to the top and say "Look this is something of tremendous importance to the country" at an early stage. I wanted to first make sure the idea was sufficiently sound, and that there was some hope of getting it developed. In fact, the interest of the Post Office rather surprised me - the fact that so many people turned up, and that they came from a fairly high level. From then on, relationships were always extremely close, but I realized after a while that I was talking to a particular group in the Post Office Engineering Department - George Allery and Phil Kelly come to mind. They were a particular group who had obviously been inspired by this idea and were very keen to develop packet-switching, and really sold on the idea. But they were a very small group, and from time to

time I got the impression that there was a very strong resistance within the Post Office to the spending of any real money. And this is interesting because the Post Office started a number of data communications services which completely floundered for lack of interest. They were convinced that if they provided this type of service everyone would want to use it. They went out and spent an enormous amount of money on promoting it and setting up equipment, and got no takers at all, and it rather annoyed me that the Post Office on its own could go out and decide what it thought the public wanted, but was not prepared to spend money on an idea I thought the public wanted. Now, in fact the four studies which they carried out were the first time they really put effort into this. I was involved in the discussions to sort out what these tasks would be, but in fact they each went away and did almost something quite different, depending on the interests and enthusiasms of people involved. But on the whole, the conclusions were pretty favourable to the idea of developing packet-switching. So ... there was no problem between me and the Post Office; I think any politicking, or any internal problems, were actually ones that took place within the Post Office.

MCK After the seminar you wrote the famous report in June 1966 entitled "Proposal for a Digital Communication Network". Who did you send the report to, and what were their reactions? Incidentally, it seems extraordinary to me that you never did publish that paper, to give the ideas the widest circulation: why not?

DWD As to who they were sent to, I don't know if we do have a record. There was a fairly large circulation list for reports to which they went automatically, and then on top of that I must have chosen all the people I could think of in the Post Office. That was the first time that I sent a report out with the aim of reaching all the people who might be informed. More than that I really can't say. Not many people gave reactions to it actually, but those that I did get from talking to people afterwards were pretty favourable. On the whole, I was beginning to get very encouraged, and to feel that the ideas were very significant and would also be acceptable to other people. So on the whole the report was well received. Now one copy of it certainly went to Larry Roberts and, when I visited him in the Pentagon on one occasion, it was lying on his desk in tatters. It had obviously been very heavily thumbed and turned over, and he grilled me on a number of aspects of it.

MCK What date was that?

DWD That I can't remember. It was in the days when the ARPA network was just beginning. I think it consisted of about three nodes, so it was very much early days for the ARPA network. My interest in it was to find out how much Larry Roberts had really gathered from what we did. The impression I got, from things he said, was that it was not so much the technical ideas in my report, but the fact that we were enthusiastic and believed it would work and that it could be made to operate quite easily. That is, I think, the main criticism that we heard at the time. It would work, but how difficult would it be first to convince people to adopt an entirely new way of communications. And secondly, all experience in message-switching showed that the software problems were very difficult. Well this is absolutely true, they are. I think the reason was that message-switches were being made to try to operate like a torn tape centre. In other words they were trying to emulate the human system, which was really quite complex and involved a certain amount of subtlety in dealing with particular kinds of failures. The idea of sending messages - packets - and relying on lack of any response to send them again was unlike the way human systems work. I believed that by doing it in a number of layers in this way - having an end to end protocol, and protocol over each link and so on - one would make software problems a lot easier, and I think it did turn out to be so. The general impression was that software problems would prove to be quite intractable. In fact, in our 1968 lectures at Edinburgh University [at the IFIP '68 Congress], one of the ATT people present said in a discussion, rather forcefully, that he thought it was a nice idea but one that would prove to be quite impracticable because of the enormous software problems. That was at the stage where we already had some experience of the software problems at NPL, so we were quite convinced we could actually make it work. But generally speaking, the response to the report was certainly not "anti", but I can't say I had a lot of response. The main difficulty was to get people to take an interest.

MCK I think you've answered question 9I [Question 9 was: In August 1966 you became Superintendent of the Division and started an NPL research project on data communications. Was this in someway related to your inability to get the Post Office to take a more lively interest?]

DWD Yes, I think when you say "the inability to get the Post Office to take a more lively interest", I didn't expect them to actually, from the start - I was amazed that they took any notice at all. I mean, I had been in contact with them enough to know they were a pretty large, monolithic, organization, in which to get anything done you have to convince a lot of departments. The fact that my ideas had had any impression on them at all was to me rather amazing, and I didn't expect them to put a lot of effort into it quickly. What I thought was that by popularizing the idea outside, by showing them in the experiment that it would work, we would eventually be able to put pressure on the Post Office to do something, and that was how it turned out. I think it would have been impossible for someone else at the Post Office, no matter how valid their ideas were, to go in and convince them at that stage to make a kind of U-turn in the type of switching they carried out. That was inconceivable. Probably what we did was about the best that could be achieved. Although I think of course it could have happened much faster. The resistance was a bit too strong.

MCK When you became Superintendent you launched the data communications research group that was led by Roger Scantlebury. He reported to Derek Barber, and in turn Derek Barber reported to you. You had many other responsibilities at the time, and you couldn't take a very direct part in the project. Well, I know you are a good research manager, so I'm interested to know how you kept control of the project. Were there regular meetings and reports? How did you keep in touch with the technical developments?

DWD Well, I have to admit that I probably had neglected my management tasks. When you say a "research manager", one of the problems in managing a large group of perhaps a hundred people is that you have all kinds of completely non-research things to worry about. I mean very mundane things. On the whole I tended to treat them rather lightly and dash off answers to serious questions without too much thought. I don't think I got a very good reputation in the NPL as a manager. My main task I was able to pass on to others. Most of the work of the Advanced Computer Technology Project was taken on by first of all Lou Page. The other projects in the division I could keep an eye on by means of formal meetings, once a week or once a fortnight. But with regard to the communications project, I took a much greater interest, talking to people virtually every day. Partly because I spent a lot of my time outside the NPL promoting the ideas, lecturing on the subject in innumerable

conferences and so on, and writing what were mainly conference papers. I think one of the questions you'd asked earlier was "Why on earth didn't that paper get published?", and I really hadn't any good reason for that. I suspect that it would have been quite difficult to get it published in a learned journal because it doesn't have much in the way of original ideas in it. Maybe if I'd gone to a lot of trouble to rewrite it I could have got it published in a prestigious journal. At that time it was certainly much easier to get papers published in conference proceedings and I took the easy way out. I think that's the trouble - I mean conference proceedings you can easily get published, but they don't make such good references for the future. Of course ARPA was very similar. ARPA papers were mainly published in conference proceedings, so we were both doing the same thing. And that particular paper would have needed complete rewriting. In fact I published half a dozen papers of a similar kind in various conference proceedings around that time. But you're quite right, I think - a bit more care in publishing that much earlier would have been valuable. Of course the right types of journals didn't exist. The ACM I think would have been very snooty about it. The only place I can think of that would have taken it in those days would have been the Transactions of the IEEE. You could get papers published within six months there, so I might have done better to have done that.

MCK By July 1967, the project seemed to be on a sound footing and you applied to the NPL Steering Committee for about £120,000 funding over the next three years. This must have been quite a large sum: did you meet any resistance, or did you have to camouflage the application up in any way?

DWD I don't remember too much about this. I don't remember having any tremendous problem in meeting resistance, so I think I must have had some quite good support from the Director. I have no memory of difficulty in getting support at the NPL. Later on, when the Requirements Boards came along, I was lucky again in that the board had just the right people on it. They were over enthusiastic about what we were doing really; they said "Yes, you must do this and why don't you do that and so on". The result was that we ended up getting their support for more people than we actually had, and since there was no way of recruiting extra people on that basis, it was no argument. You could never produce an argument for increasing the staff simply because the Requirements Board had said you had to do more than you asked for. At that time they had the idea of bringing in Harwell, and a group was started at

Harwell essentially to support the work at NPL. In fact, the collaboration didn't work that well, in that Harwell went off entirely on its own. But for most of the period from then on, we had regular meetings every few months with Harwell to talk about our research projects. I wouldn't say it was a close collaboration though; Harwell did really their own thing.

MCK In several areas or in data communications?

DWD In data communications. This all began at the time of the first Requirements Board and I can't remember when that was, a little bit later than we are talking about now. [The Computers, System and Electronics Requirement Board was established in 1973.] But I certainly don't remember having any great opposition from the [NPL] Steering Committee. Initially, I dressed up the project only before I became Superintendent. I wanted to get the work started, and at that point we did start some work by including it under the heading "Programming Research", which was one of the headings in our research program which at the time we hadn't much effort to apply to. So the initial formative work was done under that heading, but that was the only dressing up we did.

MCK In Autumn 1967 Roger Scantlebury read a paper on behalf of the group to the ACM Symposium on Operating Systems at Gatlinburg, Tennessee. At the same meeting Larry Roberts gave his first paper on the ARPA network. Did you know about the ARPA network at this stage? And what were your relations with Roberts over the next few years? You've mentioned one meeting.

DWD Well the relations were very close and cordial over the whole of this project, but the meeting in 1967 I think was probably the biggest influence we had on the way things developed.

MCK Had you heard about the ARPA network?

DWD Yes, I think we heard about it very early, but I'm not sure whether we'd heard about it before that meeting. (You might deduce that by looking at the wording of Roger Scantlebury's report of that meeting.) Now, as regard to Larry Roberts, while he was still at MIT he had written a pape: on multiple computer interaction and he did an experiment between, I think, MIT and the Lincoln Lab. (he was actually a Lincoln man). But that was quite a simple

sort of protocol which he set up to exchange messages; it didn't amount to very much more than that. I think that was in fact [a paper by Roberts] "Towards a Cooperative Network of Time-sharing Computers" in 1966. Now, at the Gatlinburg Symposium, he wrote a general discussion about the requirements for networks, and for the way in which networks should operate. He talked about ARPA and the way it was going to support a network project, but the actual communications scheme that he described was extremely simple-minded and was based on only low data rate lines. If it had used store-and-forward, which wasn't at all certain, it would have been of the message-switching kind because all the lines were going to be of low speed - about two kilohertz. The only remark about store-and-forward that he makes is right at the end of the paper where he says in his final sentence "This is very wasteful of the line and unless faster dial up times become available, message switching and concentration will be very important to network participants". And that's the only remark in the whole paper, which is surprising because, two years earlier, the discussion at NPL showed that the idea of store-and-forward was at least around. So clearly he hadn't really thought about it to anything like the extent that we had. This was the remark made by Roger Scantlebury, when he came back: in the USA the thoughts that we'd been having had really got nowhere. According to Larry Roberts, it was our enthusiasm for store-and-forward paper at the '67 meeting, and the '66 paper which I sent to him, which made them decide to go into packet switching. Some of the terminology shows, I think, that they were heavily influenced by this paper. According to Larry Roberts, anyway, that was the reason they went in that direction and from that point of course they brought in BBN. We didn't have much direct contact with them during the six months or so that they developed their IMP design, but we were extremely surprised when we saw the IMP design to discover that they used the same message packet size as us, and they had many things in common. Except that they had no interface computer: it was at first only a means of connecting mainframes and there was no access by terminals except through a mainframe as host, a sort of IBM-like solution. The idea of a terminal IMP was not exactly an afterthought - because it must have been in their minds - but it was a second stage of the project. In that sense our project, which was essentially a terminal IMP, was quite complementary to the ARPA network project. I don't think many people saw it that way, but that's how we thought about it. We had good relations with Larry Roberts through the whole thing. I rather feel that, partly because of the enormous publicity that surrounded the ARPA network

and the fact that most people in the USA had heard of nothing else, Larry Roberts gradually came to forget the early history and to believe that packet switching began with the ARPA network. That strictly isn't true, although I think it was only in the very early stages that we had much influence.

MCK While all this was going on, Rex Malik and other people founded the Real Time Club, with Stanley Gill as its Chairman. When did you first become aware of the Real Time Club, and why did you decline to join it?

DWD I was asked to speak to them in their early days. The Real Time Club hadn't taken on data communications as a kind of main theme at that stage, but I believe that the talk I gave really made them decide that they would put a lot of their political effort into promoting data communications. I went to two or three meetings after that, but I decided not to become a member because it was openly a political lobby intended to put pressure on the government. I was rather worried that I shouldn't therefore belong to it, that I should remain much more impartial. Of course if I'd asked the officials at NPL they would have said that was right, but I didn't ask them. I just on my own account decided to be a little bit cautious and not to join this highly political body, except in the background. Of course the [July 1968] meeting they held at the Festival Hall I was very heavily involved with. It was really very interesting, and it was like plotting because I wasn't supposed to be involved with political pressure, but in fact most of the planning for that took place at NPL. (A man called Dunlop was one prime movers, if I remember rightly. There was another curious character, whose name I've forgotten, who was present too. Rex Malik provided a lot of the enthusiasm as well.) We used to meet here to discuss it and plan it, but of course officially I wasn't involved. That was quite important at a political level; I doubt really whether many people were influenced at the technical and managerial level by that meeting. It was, I think, above their heads mostly, because most of them hadn't really cottoned on to the significance. Stanley Gill of course spent a lot of his time on the politics and on influencing people at high levels.

... Later I was asked to go to a meeting between Wedgwood-Benn (the Minister of Technology) and the Post Master General (John Stonehouse). First of all I went along to Wedgwood-Benn to brief him on the importance of data communications, and I felt he was absolutely first-rate at picking up ideas and expressing them in his own words far better than I could. So when we came to the meeting, he really put over the whole story I had given him that

morning quite superbly. Stonehouse, by comparison, was floundering all the time and was being held up by Merriman [J.H. Merriman, Director of Post Office Engineering.], who was his chief engineer, and who kept stepping in and helping him out. The result of this meeting was that there should be top level meetings regularly between the two. Well, the Post Office started off with top level meetings with Merriman chairing them and so on, and gradually these meetings went down and down the hierarchy, more and more junior people taking them. They still continued, and there was until only a couple of years ago a regular meeting every six months. Which was important for me because it was a point at which I got the latest information about what was going on in the Post Office, or British Telecom as it has become. But then with the departure of Phil Kelly, who was the last chairman of that meeting, they fizzled out completely. It was an interesting example of an initiative taken at a high level which by action of the various civil servants concerned was gradually deflated until it became a very formal interesting exercise but of no great value.

MCK What was the reaction of your contacts within the Post Office to the Festival Hall event and about the formation of this committee.

DWD No, I don't think they were very worried. They were very self-assured and confident that they knew how to run telecommunications, and weren't going to be influenced by anyone else. I thought that Stanley Gill's speech, which has been reprinted many times, setting up two alternative scenarios was a masterful piece of political satire, almost. And as people have pointed out since, its turned out to be almost exactly true: the pessimistic scenario was followed almost to the year. I think this resulted in the Bowden Committee, which I attended once or twice, but I don't believe any of that was very influential. It's possible, though, that it did finally nudge them into the EPSS. But what it needed really was a much more heavy political push, or push from someone inside the Post Office. This would have got the thing started really practically, like TRANSPAC which the French PT team put a lot of effort into, and overtook EPSS rapidly. EPSS also suffered by being too early, so that it used a very clumsy interface because X25 hadn't yet arrived, so that altogether it probably didn't do a lot of good. It didn't work out too well as a test on the users requirements because not too many users emerged. It didn't really solve many technical problems, and I think either a slightly later start

using X25, or a somewhat earlier start trying to rival the ARPA network, would have been much better.

By the way, you were wondering about my own contribution to the report to the [Post Office] Economic Development Committee. This was all part of a number of initiatives that were going on. I made several such proposals to different committees and to different groups. For example, Ieun Maddock started up something for a Mintech network which we started to discuss for about six months and then dropped. Much later on there was a defence network, called Grid 77, which was discussed at great length, and in great detail, which would have been a packet-switching network for all the non-operational aspects of the Ministry of Defense. It would have been a tremendous project and would have advanced the course of networking very rapidly, but in fact it ended as a fiasco because the three services couldn't possibly bare to share any computing power with the others. We had a meeting in which the Airforce was arguing that it would be absolutely disastrous if there was a computer centre working for all three services, and a bomb dropped on it. Everyone tried to point out that it would be far more disastrous if there was one centre working for the RAF and a bomb dropped on that. But the argument wasn't going to work that way. The inter-service rivalry killed Grid 77.

MCK Do you agree that Stanley Gill played a major role in the British data communications scene? I would be interested to know your opinion of Gill, at a personal level, as well as his technical and political contributions.

DWD Yes, well of course everyone knows about his early technical contributions. (Incidentally he worked at NPL and when I first worked here. He was responsible for input/output equipment of the ACE, which was punched cards mainly, and when he left to go to Cambridge I took over. So I knew him very early.) On a personal level, I found him a delightful chap to work with: very pleasant, very modest, perhaps not forceful enough for the kind of political role that he undertook. Although I must say he was willing to spend a lot of time trying to influence people at the top level, which I found to be rather an uninteresting subject - I would prefer to get on with the technology and leave that to others. So I think he tried to do something that perhaps I should have done more of. I think he did play a major role in the sense that, to the extent that people at the top level did get influenced and appraised at this, he was responsible.

But it's extremely difficult when the whole of the Post Office is thinking another way. For example, Merriman was the Chief Engineer, and took a great interest in the subject and I had several long conversations with him. Sometimes, at the end of the conversation, I was convinced that at last I'd made him understand what data communications was all about, and it wasn't just a matter of providing modems, and even that data communications might become important as an internal part of the telephone network. A theme I kept hammering away at was the control network of the telephone network, which at the time were being redeveloped. A new system was being developed at CCITT, and should in fact be based on packet-switching, and that such packet-switching could actually be shared with the public, provided you made quite sure that the essential part of the telephone network was well protected. All these are things which still haven't developed fully, and Merriman was convinced of them at the time. When I spoke to him only a month later, I found that he had completely forgotten the discussion and moved back again; in other words he was influenced by whoever was talking to him at the time. Really you can't fight against a large organization like the Post Office. Ninety per cent of the people are thinking along one particular direction; maybe someone at the top, who could be more forceful and who happened to get inspired by an idea, could have made the change, but such a person didn't emerge.

There was another aspect of packet-switching that I'd like to mention, also, which is that from the start we believed in very high speed packet-switching. In fact we talked about 1.5 Mbit/s lines, which haven't yet come into operation. I still believe that the technique has this capability of working at much higher speeds. If you look at the need for making a switch at that speed, clearly it's got to be rather a special purpose switch. When we began, when minicomputers were the cheapest thing you could buy, it was rather hard to see how that could happen. But now that every telephone set has a microprocessor in it and every switch can have dozens of processors, it would be quite easy to develop a specialized packet-switch. Now, in fact, in Toronto University at one time, one of their engineering staff on his own developed a design for a specialized packet-switch which had a very high data rate. So I'm quite sure people tended to speak about packet-switching as though it was essentially a low speed system - that if you can only have 64 Kbit/s trunk lines then users will clearly not be able to get anything like that. But in fact I don't think there's anything inherent in that; it's simply

that we've always made the switches out of standard computers and we haven't tried to design a specialized switch for that purpose.

MCK A data communications engine?

DWD Yes, it would have to be designed specially for the purpose, but no-one has really tackled that, so I think there's a lot of potential. It's remarkable that's it's taken so long for that penny to drop. What is amazing to me in telecommunications is how old ideas stick, like a telex network which is so out of line with modern technology that it's incredible, and yet it's holding out against teletext and so on. It's the enormous capital investment. Very early on, I discovered that one of the realities of the telecommunications network is that once something has got established its hard to shift it.

MCK Well, the work in NPL was moving forwards. In January 1970 the Mark I network, as it was later called, first operated successfully, and it went live in July 1971. By then of course the NPL work had lost the limelight to the ARPA network. How great was the sense of achievement in the division when the network did become a reality, or was it overshadowed by the much more widely known ARPA network?

DWD I don't think we ever felt overshadowed by the ARPA network. We were working in a slightly different area of the network and we felt, as I said earlier, that our work was complementary. Also, when we lectured in USA we found that everyone understood what we were doing. We talked quite a lot about our simulation work. So we didn't feel overshadowed.

MCK Did you mean complementary because it was a local area network?

DWD Yes, yes. Because we were handling the interface problems which the ARPA network in the early days (up to about 1972 or '73) had not even really begun to tackle. So we felt we had a very important thing to do, which was to develop the protocols appropriate to the interfacing of terminals of many types into the network. Now, of course, the actual operation of the network was fairly gradual, so there was no one point at which everybody cheered and said "Well, it's all working". There was, I think, a great sense of achievement among the people who were building this local network in what they had done - as a demonstration, really, of the packet-switching principle.

As far as the users at NPL were concerned, like most users, they're not interested if something doesn't exist. You can talk to them about some project that you're engaged in and say "Look, would you use it?" and they say "I might, I might not", but as soon as it's there they adopt it with enthusiasm. There was a lot of enthusiasm about it among the people who built it, but the users took it as quite part of the normal landscape and very quickly were using it every day without thinking there was anything unusual about it. Of course in its first form it lacked a packet interface; it was just a byte interface which would produce a communication link between two terminals, but it wouldn't enable one computer to talk to a lot of terminals by interleaved packets. So it didn't really, in its Mark I form, actually demonstrate very much of what we wanted to demonstrate. Also the Mark I network's operating system proved to be a little unwieldy, or a little inefficient. The Mark II therefore had a very much improved operating system, and worked much faster, and also had the two types of interface.

MCK It was part of your original proposal to have a packet interface. What had happened that the Mark I system developed as a character switch rather than a packet switch?

DWD ... we wanted to simplify the thing and get it going easily and quickly. I was a little bit against that. I think it was a pity, really, but it was far better to have something demonstrable, that people would actually use. By that time, people were clamouring for its use. We had already got cables around the place, and were using the NPL standard interface around the laboratory, so people were really clamouring for it - at least a few of them were. So I think the Mark II network was clearly what we were really aiming at, but we were probably right to put together a rather simpler system in the first place.

MCK This wasn't something that suddenly occurred to you when the project had finished; that "Oh dear, we should have had a PAD in there."

DWD No, no, no. It was an essential part: in the earliest papers, the PAD was there. I think maybe, in retrospect, it would have been better to have built the thing in the Mark II form straight away. The difficulty, from my point of view managing the thing, was the way the software was designed. They thought about it very very hard before they designed it, and only then did they begin to write it (which I understood was quite the correct way of doing it).

But the result was that a year had passed during which nothing had been demonstrated working at all. I mean the elementary hardware was working but the network as a whole couldn't be demonstrated and I was getting quite worried that perhaps the whole thing would fold. So when they said "OK, we can put something together much more quickly if we scrap this part of the thing and if we do that as a second stage", the idea of seeing something working was quite attractive. It had been a long time during which nothing had actually been demonstrable, so its arguable both ways. Anyway the Mark II network was what we were really after and was really the intention the whole time.

MCK You've indicated that you were disappointed by the performance of the Mark I network which was I think 150 packets per second throughput. Although that was much worse than your early optimistic forecasts of a couple of thousand, it was actually very much in line with your simulation work and studies going back to 1967. So I wonder why you were disappointed when it was within a factor of two comfortably.

DWD Yes, the disappointment was spread over some time. When I say I was disappointed, I mean that strategically it was a disappointment because I had hoped that we could demonstrate something much closer to the original very optimistic thing. But the simulation, as you say, had already shown that that wasn't going to happen, so the disappointment was rather earlier really. I think I had simply not understood the overheads that you get in software.

MCK Another aspect to the work at NPL, which was your collaboration with Wyn Price on flow control in networks and congestion avoidance. Can you tell me a bit about that work? What was the contribution of Wyn Price and what was your contribution?

DWD I was intending, right from the beginning, when I had my first discussion with Derek Barber about what we should do about this new idea, it was to be a two-pronged attack: an experimental network and a simulation. But it was quite difficult to find the right people, and of course we had to find someone who had the right knowledge for the simulation work, and who had been freed by another project. Roger Healey started this, and then when he left I think [A. R.] Meetham did a bit of work for a while, and then Wyn Price took over. Wyn Price made it much more of a kind of a career than the others had done.

Now, initially, I would pose a problem and say "Look, this is how we might simulate it", and they would then go away and actually write the program. So I was very much in control. Initially I would say to Wyn, "Why don't we test this, or try that?", but after a while of course, Wyn took over and was proposing his own experiments, and doing his own thing. Of course, I'd forgotten, Costas Solomonides came in and helped Wyn for a while, too. I think at that point, it became much more of a self-generating project in every way. But initially, I had been enthused by simulation studies because when the ACE Pilot model first worked, I decided that that was one of the things I was going to use it for. I did some early simulation work on safety in mines, and also on road traffic control - very very early work, the first work ever I think, on simulating road traffic control. So I had a background in this, and wanted to do it, but it was in fact Healey, and then Meetham and Wyn, who made it possible. Though Wyn's work, as I said, gradually took over and became self-generating. I think you wanted to ask about the permits...

MCK Yes, what was the importance of the isarithmic concept? Tell me a little bit about how the concept arose.

DWD It arose really from Healey's work, which in a number of diagrams seemed to me to indicate that the throughput of a packet-switching network would rise as the number of packets in the network increased. Clearly, you had to have more packets in the network in order to get more throughput, and then beyond a certain point it began to fall. This was a kind of instability which eventually meant that if you loaded it with too many packets, none of them would move at all.

MCK That's very much like a time-sharing system, isn't it?

DWD Yes. That curve, which I remember as just a bump which went back to zero, impressed me very much and it seemed to me that we should therefore try and control the number of packets to be never greater than the half-way point. Then we would avoid this type of destructive congestion. As a means of controlling the number of packets, it occurred to me to give them a permit before they could get into the network. The idea was, I think, a good one, and in simulation it seemed to work, but it had the danger that permits could get lost and so the thing had a nasty collapsing process if some part of the

network was doing the wrong thing. A Frenchman developed a scheme in which permits were continually being created and destroyed which would be self-healing in this respect, except that I think if you bled too many permits it would probably not work too well. I don't think the idea was ever more than a theoretical one. I published a paper on it because it seemed to me an idea worth talking about, and it was quite interesting to statisticians who liked the idea of this, and people generally found it an interesting idea. It was also my conference paper to get me to Ljubljana! I never thought of it as much more than that; I didn't think it was going to be a practical issue I'm afraid. We called it the isarithmic method. The word is strictly derived from greek. I went to a greek scholar in the division, and asked him for a simple name and we worked out "isarithmic".

MCK What is the current significance of isarithmic congestion control? Is it the case that networks are not on a large enough scale yet to have really brought other techniques to the fore?

DWD Yes, it's a very interesting question. Right from the beginning by doing simulation we found all these snags. And ARPA network found some interesting protocol lock-ups, not just by simulation, but by actually locking the network up. The same types of congestion can occur in circuit-switched networks but for some reason it had never been realized. When the Alaska earthquake took place, everybody in the USA, was phoning up to find out what had happened. There was a rather clever adaptive routing scheme in the USA, which meant that if you couldn't get through it would try other routes, even to the extent of going across the country and back again, just to get from Washington to New York. They had a very elaborate adaptive routing scheme, which is actually one of the mechanisms for making this type of destructive congestion, and the whole network seized up at the time of the Alaska earthquake. There was a big inquest on this, and papers were written, which when we read them, sounded exactly like what we had been doing with packet-switching. They were discovering exactly the same kinds of rules of thumb to make a network more stable, and what we realized was that in developing a brand new network we had gone much further in understanding its possible maladies than had actually been done for the telephone network. Now the existing packet-switch networks, based on virtual circuit-switching, of course don't have this kind of type of congestion problem in quite the same way. The congestion problem is solved, in my view, in a rather cruder way.

Nevertheless, if one wanted to look at their congestion problems, it would be connected with setting up new circuits rather than sending packets over existing circuits. So, I think, what we did was rather more appropriate to datagram networks than to today's packet-switched networks. But it is interesting that we went further than had been done with the existing networks.

MCK In 1973, you published a book with Barber (and with the help of Price and Wilkinson) "Communication Networks for Computers". Sometimes a book can be more influential than the system it actually describes. How much of an impact did the book make? And did you get any kind of feedback from the people who read that book?

DWD You might be interested in the reasons for that book. In 1969, I was asked by the Japanese equivalent of NPL in Tokyo [the Electro-Technical Laboratory] to give some lectures. They were paying the expenses and so on, and in return for giving my wife and myself a bit of a holiday, I agreed to give, I think it was seven three-hour lectures. You have to speak rather slowly, so it was not quite the same as seven three-hours here, but it was an awful lot of lectures, and I spent a lot of time preparing material. The interesting point about the Japanese lectures was that everybody there seemed to have read our papers in great detail. They all had copies of them, and once the discussions started - not during lectures, but a quite separate discussion period - they were clearly extremely interested. I was struck by the fact that interest in Japan was in that sense far greater than if you'd set up a course for about twenty people in Britain. They were obviously on to a good thing quite early. Because we had all that material, I decided it would be a good idea to put it into a book. That was the origin of the book. It was actually altered a great deal for the book. I think it was influential, in the sense that it was read by lots of people. I know this because I'm continually meeting people who say "Ah, you're the Davies who wrote that book". That book, and the one that followed it, have been read by a lot of people. I think it's sold about eleven thousand to date, and the second book about fifteen thousand, which is quite a large sale. So, yes, it's influential in that sense. I often meet people who clearly learnt about packet-switching from that book. More than that I don't know. I would imagine that publicity about the ARPA network and the success of TRANSPAC and so on, had more influence on the PTTs in other countries than just reading a book. When they see something actually working it is much more

impressive, so I wouldn't over estimate it's influence, but it certainly reached a lot of people.

MCK Which was the better book?

DWD Oh, well, the second one nowadays. The first was more original in that there was no other book like it at the time. The second book, however, was better written and brought the subject much more up to date. Although of course that is out of date, in the sense that OSI hadn't been invented at the time. But I don't think I can answer the question.

MCK After about 1975 the data communications work at NPL began to wind down, with several of the staff moving on to new projects. You retired from being Superintendent in 1977 and became a "special merit" scientist at NPL. You seem to have become an internationally known expert in computer security. It all sounds like the wheel turned full circle, because that was how you first started on data communications. Tell me a little bit how you got interested in cryptography and data security.

DWD Well, yes, I took the usual childhood interest in cryptography, I think, that many people do. But quite early, about I think, 1963 or it might have been earlier than that, we were asked by Midland Bank to test the security of ATMs which used cryptography and had been built by a small company called Speytech. It was one of the first three types of ATMs to go into use, not with magnetic striped cards of the ordinary kind, though the Midland did have a magnetic stripe on it, and at the time the other banks had punched cards and things of that sort. Well, we broke the thing immediately, and we gave them some advice on how to improve it. They came back with a cypher built into it which was a little bit more complicated, and it took me a whole weekend to break that particular cypher. So that was my first introduction to data security in a practical aspect. From that point on, banks did come to us occasionally for advice in the area of security, but usually quite short-term things, of no great significance. I mean, there was nothing to convince me that there was long-term an interesting problem for banking in security of its communications. Now, during that time I kept up an interest in cryptography, and in fact when the National Bureau of Standards sent out its request for proposals which led to the data encryption algorithm, I put in a proposal, which of course got nowhere because none of the first requests got

anywhere. I now see the proposal wasn't very original or very good; it was rather ill-defined. But, I can't quite remember why, I had been in touch with GCHQ about data security and realized that it would be very difficult for NPL to get involved in this area, because of the way they regarded any government work in this field to be very much their province. So I had several discussions with them, including the discussion about the proposal I put into NBS. So somehow I was generally aware of, and interested in this field, but not working professionally in it. And then in 1976, the data encryption standard began to be discussed, but I had already seen it earlier, say around about 1975. I had already seen the work at IBM, the Lucifer, working, in 1972 and discussed it with the inventors; ... and in 1976 and 1978 work on the public key systems came out. So at that point I was sure that the expected increase in interest, the thing we had been anticipating all along, was actually beginning. There was now enough technology to make it a worthwhile subject, and I should move into that area. I'd become a little bit disenchanted with data communications because it seemed to me that most of what was going to be done in the future outside the PTTs was going to be involved with the honing of standard protocols and getting them exactly right. I wasn't that interested in the fine details of getting protocols working, so I decided to make a move, and it coincided with the opportunity to give up being the Superintendent. So if you like, my interest in cryptography, like my interest in data communications had been going on for a long time. I just saw the opportunity of making the move in 1977.

MCK Well, a couple of general questions to finish off with. It is easy to be wise after the event, but nobody gets everything right the first time. If you had your time over, are there things that you would do differently? First of all in the NPL project, and secondly in the UK national data communications scene.

DWD Well, I think in the NPL project our hands were more or less tied, as we had only a limited amount of resources. I think we could have done better in the NPL network if we'd made the hardware design a little less esoteric. It came about, as you probably know, because it was based originally on an idea of Plessey's, for multiplexing. We took that over; it was a nice idea we thought, but it's a very esoteric kind of design which is hardly likely to appeal to anyone else.

MCK An alternative to that would be a conventional PAD?

DWD Yes, I'm not sure what the alternative would have been. Quite early on, when the ALOAHA was being developed in Hawaii, which led to the Ethernet, we began to work here on a kind of radio Ethernet. We would use transmission by a leaky cable, so that you would have mobile terminals that would be working like Ethernet. I think this would have been a very valuable piece of work, but we were discouraged by the rules and regulations concerning transmitting and receiving information. We tried to get a license to use some spare television bands, but we were turned down and we never fought the thing very hard, and so we never started on that project. That was a pity. I think we should have decided to stick more firmly to local area networks and continue to innovate in that field, and then we would have been in the mainstream of local area networks. As it was, we'd finished our work on local area networks before anyone else started. So hardly anybody realizes that we did work in local area networks, and we're not part of the current craze for LAN technology. So that's another thing I think we might have done: to stick more firmly to local area networks for a longer period. As far as the national scene is concerned...

MCK You've already mentioned one; that was that the EPSS probably started too early.

DWD Yes, that was almost an accident. We were keen to start it as early as we could. But having failed to start when we did, it would have been quite clever to have said "Look, by waiting another year we can hang on to an international standard". I was aware of the international standard, and I took part in the earliest discussions at CCITT, and I was present at the meetings when the X25 began. But at that point we had been pushing for so long, I think we would have been extremely sophisticated if we had realized that by delaying we could have improved the value of the experimental network. But it would have been better if we'd done it that way. As far as political pressure and so on to get things done, I suppose that I should have taken less interest in the technology and spent more time on the political aspects. I probably could have done that, because we were involved with junior ministers and so on, in the discussions on the changes that took place when the Post Office first became [British] Telecom. I think we could have had more influence at that point, if I had spent more time on the politics. In some ways the director of NPL warned me off

that, because I had at that time quite a powerful voice in the form of the junior minister. I used to advise him quite a bit, but I was warned by the director not to get involved too much with political arguments, partly because ministers are very temporary and inclined to disappear, so that all the work is wasted. And also because it's not the sort of thing for a civil servant to do. So it's difficult to know what I could have done better. But I might have spent more effort on it.

MCK Everyone I have spoken to regards this period, 1965-75, as being a golden decade in the Computer Science Division. Somewhere along the line, and I suspect it may have been tied up with the attitude to long term research, things began to go down hill. Do you have any comments on that?

DWD Well, I'm not sure I know what you mean by going down hill. But I think there has been a change, largely imposed from above which is really because the Requirements Board mechanism, and philosophy behind that, has really taken a grip on the work in the division. When we started, we still had a lot of autonomy in the kinds of research we did. We had to convince a local board, but it was one that we had some control over the formulation of. The idea of the Requirements Board developed into the philosophy that they should be largely run by industry, and that industry knew what was needed in the way of research, including all research of NPL. Now that I never subscribed to, because of course industry certainly never told us to work in data communications, nor did it tell us to go and work on computers. Industry, including IBM, was way behind in the development of computers. It was a long time before IBM had any knowledge that computers existed, and were being built in universities and government departments. So I think the general philosophy that you have to rely on industry to tell you what the real needs are is actually a misunderstanding. I guess this was long term research because it was roughly ten years from the initial inception (well about eight years) to the point where it became widely understood. About 1973, I mark as the turning point, where if I got up to talk about packet-switching I didn't have to begin by justifying it in any way. I could simply assume that people knew that it existed, and it was a good thing, and talk about it. So I think maybe eight years to get accepted, and all of ten years to become a working tool. That to my mind is medium-term research. I don't think you can do anything that's fundamentally new without a ten year time-scale. This is even true of data security: apart from banks, the need for data security and

the techniques of data security are still largely unknown in this country. So I think any worthwhile research is going to be of the order of an eight to ten year time-scale. Industry may have some very good ideas, but making everything dependent on a board which is dominated by industry I think was a great mistake. Now the other effect of this has been that committees have tended to devise the whole research project, and they tend to have a horror of duplication. They urge to people to collaborate and amalgamate into joint projects, as you can see in Alvey and so on. Now this is fine, up to a point, but if too many people get involved it becomes extremely unwieldy and I believe that Eureka is proving to be even more unwieldy than Alvey

MCK Do you think this project could get off the ground nowadays, under Alvey or whatever?

DWD It might. Because Alvey is being driven by a few people, it might just have a chance. But under the Requirements Board mechanism it wouldn't, because they would have immediately said " This is not your job, it's that of BT, who know all about telecommunications and we must not even begin to think about it". But in the event, many of the projects now involve quite massive collaborations. For example, you must have NCC in because they are going to be the representatives of the users, you must have a manufacturer involved because they are going to be representative of the supplier, and NPL therefore takes a part in a joint project with lots of other people. This takes an enormous amount of peoples' time, particularly at the top level, sitting on committees, arguing for funds and for their part in the work and so on, and becoming very political. It also reduces greatly the enthusiasm of the individual researchers to know that they are just one cog in a large system. So I believe that to that extent things, as you've said, have begun to go down hill. This is really to some extent because of pressures imposed on us from above. I think, of course, there's always a need to organize and control research because it can be very expensive, but it's also quite easy to destroy it's long term value by too heavy control. Much of what I have seen in Esprit has been studies of what might be done, or very abstract studies leading nowhere. I've done one myself for Esprit, which I could see from the beginning was not going to have any great practical outcome, because of the way it was formulated. It was really as though they were frigh.ened of real research, but rather concerned to study what was needed all the time. I've seen projects which could have been made to work within a few years,

resolved in a series of feasibility studies, requirements studies, and so on, one after the other, going on for ten years with still no sign of anyone actually making it work. So I have somewhat of a horror of the way in which research is now run by high level committees. That I would think is something of a symptom of the lack of what you have called a golden period!

MCK Well, the last question, and a technical one. Now, packet-switching arose because in the mid-1960s it took 20-30 seconds to establish a telephone call. It was therefore necessary to hold a line for the duration of a session of a computer interaction. Packet-switching offered a way round this problem by "time-sharing" a line between several subscribers, in a way analogous to a time-shared computer system. Colleagues, who know much more about this than I do, assure me that in the not too distant future we can expect very rapid circuit switching to become possible. When that happens, do you think that packet-switching will be superseded, or do you think it is here to stay for the foreseeable future?

DWD Yes, this is of course a question that has been posed many times during the history of packet-switching, and I think we are closer now to having very fast circuit switching than we have ever been. The difficulty is one, really, of economics. The way circuit switching is handled by the telecommunications people (who are of course dominated by the enormous telephone market because that pays their salaries), is to combine the two so you have in ISDN a 64 Kilobit/s switch, which is going to handle both speech and data, and therefore be much more economical for the switching of high volumes of data. That's fine, yes. As for working very fast, the economics of the switch will always be influenced heavily by the requirements of the telephone network, which could certainly improve on 20 to 30 seconds, but it doesn't become very worthwhile to get down below say 5 seconds, which most people would regard as a very good telephone response. Maybe switches will do better than five seconds, but there will never be the economic incentive to go down to very fast circuit switching. But even supposing you could get down to, say, a tenth of a second, they still wouldn't handle very effectively some of the fastest requirements for interaction. But at five seconds, it would be quite out of the question for say the type of fast interaction you have in an ATM or an EFTPOS terminal. So I think that packet-switching will still be needed for some applications, anyway, and fast circuit switching wouldn't be physically capable at a reasonable cost of doing it. Another point is that

packet-switching provides a different kind of interface, and the X25 interface-complex as it is - is much more suitable for a service which is handling a large number of terminals than having fast circuits being set up by each terminal whenever it has anything to say. So, I think the X25 interface, even if it's backed-up by something different from today's packet-switches, is here to stay. And again the inertia factor comes in: once you've got a lot of people dependent on this, and you've got ISO protocols built around it, then it's going to be very hard to change it.

Now, several things can happen: one is that packet-switching can go very much higher in speed and I think that will happen for certain applications; another thing is that fast circuit switching could come in and handle certain applications, though I don't really see that as being a likely development, for reasons I mentioned earlier. If packet-switching exists, and circuit switching in the region of one or two seconds can handle a lot of other requirements for wide band data - like facsimile, low quality video, and things where an odd second doesn't really matter - then there's no economic incentive for very fast circuit switching. I think many people look at data communications without thinking too much about the intelligent interface problem. In the early days, on the same occasion when I lectured to CCITT on packet-switching (I think it was 1969 or 70), an American from IBM [Jack? Warden] said "What we need is high bandwidths, high data rates; you give us that and that's all we really need". But if you can imagine that everybody can be given a 64 Kilobit/s data rate from their terminal into the central computer, now you've got a hundred terminals coming in, and a hundred times 64 kilobits is being flung at you in time-division multiplexing. What do you do with it? It's an impossible task to sort that out. The first thing you have to do is unscramble the time-division multiplexing, and virtually putting it into a packet-switch in order to get it into the form in which you can use it. Now it doesn't make sense put all the complexity into the host instead of spreading it through the network. So I think really there will be a need for both types, and I think that the packet-switching principle will probably be used to a greater extent inside the mechanisms of the telecommunications network. If you look at some of the proposals for the ISDN, you find that the whole of the control of ISDN, from the user inwards, is being done by packet-switching. So although ISDN will happen, because you have to find some way of finding the best use of the very expensive line between the terminal and the nearest exchange, I think it will carry a number of services of which access to a PSS will be one. So I think it's here

to stay, though exactly what it's future will be and how significant it will be I don't know. I think the biggest potential user at the moment is funds transfer (EFTPOS). Large countries like Canada and Australia, with a fair spread geographically, virtually have to use packet-switching for EFTPOS terminals. Even though the argument for it is not very strong, because you could do it other ways, I think that it may become a standard thing to use packet-switching for funds transfer. As you know it's being used for CHAPS.

One of the other problems I think, when you compare this country with France, has been the willingness of the French to set their tarrifs in a way which is appropriate to the future state of the network. Right from the beginning they set the tariffs on the basis that they wouldn't be making money for quite a number of years. I think it initially it was five years and they even put up beyond that. Now in this country the tariffs have been not so flexible. For example, I still don't know whether there's a bulk tariff of the same kind that they have in France. And consequently some of the big customers, like the universities' networks, which could have been handled by PSS were lost, simply because BT or it's predecessors were not willing to develop special tariffs just to try and capture people. And they won't do predatory pricing to capture the market, so the usage of PSS until recently has still been rather poor, but it is beginning to take off now.

MCK Thanks very much indeed.